arXiv:1606.02673v4 [math.RT] 14 Oct 2016

GENERALIZED REPRESENTATION STABILITY AND FI_d-MODULES

ERIC RAMOS

Abstract. In this note we consider the complex representation theory of FI_d, a natural generalization of the category FI of finite sets and injections. We prove that finitely generated FI_d-modules exhibit behaviors in the spirit of Church-Farb representation stability theory, generalizing a theorem of Church, Ellenberg, and Farb which connects finite generation of FI-modules to representation stability.

1. Introduction

Let d be a fixed positive integer. We write FI_d denote the category whose objects are the sets $[n] = \{1, \dots, n\}$ and whose maps are pairs $(f, g) : [n] \to [m]$ of an injection $f : [n] \to [m]$ along with a d-coloring of the compliment of the image of f. Composition in this category is defined in the most natural way (see Definition 2.10). If k is a commutative ring, then an FI_d**-module** over k is a functor $V : \mathrm{FI} \to \mathrm{Mod}_k$. For any object $[n]$ of FI_d, one notes that $\mathrm{End}_{\mathrm{FI}_d}([n]) = \mathfrak{S}_n$. It follows from this that each of the k-modules $V([n])$ is, in fact, a $k[\mathfrak{S}_n]$-module.

One immediately observes that if $d = 1$, then FI_1 is naturally equivalent to the category FI of finite sets and injections. FI-modules were first discussed by Church, Ellenberg, and Farb in [CEF] due to their connection with Church-Farb representation stability [CF]. Since the paper of Church, Ellenberg, and Farb, FI-modules have been observed to be applicable in a wide range of subjects (see, for example, [CEFN] [CE] [GL] [W] [N]). The goal of this paper is to consider the relationship between FI_d-modules and a kind of generalized representation stability.

We will write write V_n to denote the value of the functor V on the set $[n]$. The **elements** of the module V are the members of the set $\sqcup_n V_n$. We say that an FI_d-module V is **finitely generated** if there is a finite set of elements which generate the entire module (see Definition 2.13). The following theorem is one of the main results of [CEF], relating finite generation of FI-modules to representation stability.

Theorem ([CEF], Theorem 1.13)**.** *Let V be an* FI*-module over a field k of characteristic 0, and let $\phi_n : V_n \to V_{n+1}$ denote the map induced by the standard inclusion $[n] \hookrightarrow [n+1]$. Then V is finitely generated if and only if V_n is finite dimensional for all n, and for all $n \gg 0$*

1. *ϕ_n is injective;*
2. *$\phi_n(V_n)$ spans V_{n+1} as an $\mathfrak{S}_{n+1}$-representation;*
3. *there a decomposition:*

$$V_n \cong \bigoplus_{\lambda} c_\lambda S(\lambda)_n$$

where the sum is over partitions $\lambda \vdash m$ with $m \leq n$, and the coefficients c_λ do not depend on n.

If $\lambda = (\lambda_1, \dots, \lambda_h)$ is a partition of size $|\lambda| := \sum_i \lambda_i$, and $n \geq |\lambda| + \lambda_1$ is some positive integer, then $S(\lambda)_n$ is, by definition, the irreducible complex $\mathfrak{S}_n$-representation associated to the partition

$$\lambda[n] := (n - |\lambda|, \lambda_1, \dots, \lambda_h).$$

One important take away from this theorem is that the growth of the modules V_n eventually becomes predictable, and this prediction is based on the combinatorial data of the partitions λ.

We will find that more general FI_d-modules are far less rigid in what irreducible representations are allowed to appear as summands of their constituent modules. To make this precise, we will first need to generalize

The author was supported by NSF grant DMS-1502553.

the representations $S(\lambda)_n$. If $n_1 \geq \ldots \geq n_r \geq |\lambda| + \lambda_1$ is a sequence of positive integers, then we define $S(\lambda)_{n_1,\ldots,n_r}$ to be the the irreducible representation of $\mathfrak{S}_{(\sum_i n_i)-(r-1)|\lambda|}$ associated to the partition

$$\lambda[n_1, \ldots, n_r] := (n_1 - |\lambda|, \ldots, n_r - |\lambda|, \lambda_1, \ldots, \lambda_h).$$

The majority of this paper will be working towards the proof of the following theorem.

Theorem A. *Let V be an FI_d-module over a field k of characteristic 0, and write $\phi_n^i : V_n \to V_{n+1}$ for the map induced by the pair of the standard inclusion $[n] \hookrightarrow [n+1]$ and the color i. Then V is finitely generated if and only if V_n is finite dimensional for all $n \geq 0$, and for all $n \gg 0$:*

1. *$\cap_i \ker \phi_n^i = \{0\}$;*
2. *$\sum_i \phi_n^i(V_n)$ spans V_{n+1} as an $\mathfrak{S}_{n+1}$-representation;*
3. *for any partition λ, and any integers $n_1 \geq \ldots \geq n_d \geq |\lambda| + \lambda_1$, let $c_{\lambda,n_1,\ldots,n_d}$ be the multiplicity of $S(\lambda)_{n_1,\ldots,n_d}$ in $V_{\sum_i n_i-(d-1)|\lambda|}$. Then the quantity $c_{\lambda,n_1+l,\ldots,n_d+l}$ is independent of l for $l \gg 0$.*

Perhaps the most intriguing condition in the above theorem is 3. It can be thought of as saying that the irreducible representations which appear as summands of the constituent modules of V, corresponding to partitions with at least d-rows, eventually appear with predictable multiplicity. If $d > 1$, then the above theorem does not say anything about multiplicities of irreducible representations corresponding to smaller partitions. For example, if λ is a partition, and $n_1 \geq \ldots \geq n_r$ are positive integers with $r \leq d$, it is natural to ask how the quantity $c_{\lambda,n_1+l,\ldots,n_r+l}$ depends on l. While we do not answer this question in general, we do find an answer in one notable case.

Theorem B. *Let V be a finitely generated FI_d-module, let λ be a partition of some integer m, and let $c_{\lambda,n}$ denote the multiplicity of $S(\lambda)_n$ in V_n. Then there exists a polynomial $p(x) \in \mathbb{Q}[x]$ of degree $\leq d-1$ such that for all $n \gg 0$, $c_{\lambda,n} = p(n)$.*

Note that if $d \geq 2$ then one can construct examples where the polynomial of the above theorem is non-constant. This is already a departure from the case of FI-modules, as one observes from the third condition in the theorem of Church, Ellenberg, and Farb.

Acknowledgments

The author would like to give thanks to Jordan Ellenberg and Rohit Nagpal for many useful conversations during the writing of this paper. The author would also like to give special thanks to Steven Sam and Andrew Snowden for informing him of their result, Theorem 3.9. Very special thanks should also be given to Steven Sam, who first observed most of the ideas used in Section 3.5. Finally, the author would like to acknowledge the generous support of the National Science Foundation, through NSF grant DMS-1502553.

2. Basic Definitions and Notation

2.1. **The Representation Theory of the Symmetric Groups.** We begin with a brief review of the complex representation theory of the symmetric groups. For the remainder of this section, we assume that k is a field of characteristic 0. A reference for much of the material which appears in this section is [CST]. It is a well known fact that the complex irreducible representations of $\mathfrak{S}_n$ can be defined over $\mathbb{Q}$. For this reason, all of what follows can be done over the field k.

Definition 2.1. For any positive integer n, a **partition** of n, denoted $\lambda \vdash n$, is a tuple $\lambda = (\lambda_1, \ldots, \lambda_h)$ of positive integers such that $\lambda_i \geq \lambda_{i+1}$ for each i and $\sum_i \lambda_i = n$. Similarly, a **composition** of n is a tuple $(a_1, \ldots, a_h)$ of non-negative integers such that $\sum_i a_i = n$. Observe that the only differences between these two concepts is that partitions are ordered, and compositions allow for zero entries. If $\lambda = (\lambda_1, \ldots, \lambda_h)$ is a partition, or a composition, of n then we use $l(\lambda) = h$ to denote its **length**. We will also use $|\lambda| = n$ to denote the **size** of λ.

Theorem 2.2. *There is a one to one correspondence between partitions of n and irreducible representations of the symmetric group $\mathfrak{S}_n$.*

Given a partition $\lambda \vdash n$, we will use S^λ to denote the irreducible representation of $\mathfrak{S}_n$ associated to λ. By convention, the partition (n) will correspond to the trivial representation of $\mathfrak{S}_n$.

The correspondence between irreducible representations of $\mathfrak{S}_n$ and partitions of n implies many strong connections between the combinatorics of a partition, and the algebra of the associated irreducible representation. Many of these connections are stated in terms of Young tableau.

Definition 2.3. Given a partition $\lambda = (\lambda_1, \ldots, \lambda_h) \vdash n$, we visualize λ as a left justified diagram comprised of rows of boxes of equal size, such that row i has precisely λ_i boxes. Such a diagram is known as the **Young tableaux** associated to λ. The box in position (i, j) is defined to be that which is $i-1$-rows down, and $j-1$-columns to the right, of the box in the top left position. If λ is a partition whose associated tableaux has a box in position (i, j), then we write $(i, j) \in \lambda$.

A **filling** of the Young tableaux associated to λ is a bijection between the boxes of the tableaux and the set $[n] = \{1, \ldots, n\}$. Any filling for which the numbers are increasing down every column and row is called **standard**.

Theorem 2.4. *For a partition $\lambda \vdash n$, write Tab(λ) to denote the set of standard fillings of λ. Then,*

$$\dim_k S^\lambda = |\operatorname{Tab}(\lambda)|$$

One approach to computing $|\operatorname{Tab}(\lambda)|$ is through the hook length formula.

Definition 2.5. Let λ be a partition, and assume that $(i, j) \in \lambda$. Then we define the **hook** at (i, j) to be the sub-diagram consisting of this box, as well as all boxes $(k, l) \in \lambda$ such that either $i = k$ and $l \geq j$ or $l = j$ and $l \geq i$. We define the **length** of this hook, denoted $H(i, j)$, to be the total number of boxes it contains.

Theorem 2.6 (The Hook Length Formula)**.** *Let $\lambda \vdash n$. Then the dimension of the irreducible representation S^λ is given by the formula*

$$\dim_k S^\lambda = \frac{n!}{\prod_{(i,j)\in\lambda} H(i, j)}.$$

Definition 2.7. Two partitions $\lambda = (\lambda_1, \ldots, \lambda_h) \vdash n$ and $\mu = (\mu_1, \ldots, \mu_l) \vdash m$ are related $\mu \leq \lambda$ if and only if $l \leq h$ and $\mu_i \leq \lambda_i$ for all i. This is equivalent to requiring that the tableaux of μ fit inside the tableaux of λ.

For the purposes of this work, we will need a version of Pieri's rule which is slightly more general than that which is usually encountered. If $a = (a_1, \ldots, a_h)$ is a composition of n we will use $\mathfrak{S}_a$ to denote the subgroup $\mathfrak{S}_{a_1} \times \mathfrak{S}_{a_2} \times \ldots \times \mathfrak{S}_{a_h} \leq \mathfrak{S}_n$. The following theorem follows from the usual Pieri's rule, as well as a simple induction argument.

Theorem 2.8 (Pieri's rule)**.** *Let $\mu = (\mu_1, \ldots, \mu_l) \vdash m$, and let $a = (a_1, \ldots, a_h)$ be a composition of $n - m$. Then*

$$\operatorname{Ind}_{\mathfrak{S}_m \times \mathfrak{S}_a}^{\mathfrak{S}_n} S^\mu \boxtimes k = \bigoplus S^\lambda$$

where the sum is over chains of the form

$$\mu = \mu^{(0)} \leq \mu^{(1)} \leq \ldots \leq \mu^{(h-1)} \leq \mu^{(h)} = \lambda$$

such that $\mu^{(i)}$ is obtained from $\mu^{(i-1)}$ by distributing a_i boxes to distinct columns.

Remark 2.9. Note that it is not always the case that the right hand side of Pieri's rule is multiplicity free. Indeed, the multiplicity of S^λ in $\mathrm{Ind}_{\mathfrak{S}_m\times\mathfrak{S}_a}^{\mathfrak{S}_n} S^\mu \boxtimes k$ is precisely the number of chains between μ and λ satisfying the conditions stated in the theorem.

2.2. FI_d**-Modules.** Fix a positive integer d for the remainder of this paper.

Definition 2.10. The category FI_d is defined as follows. Objects of FI_d are finite sets, while morphisms in FI_d are pairs (f, g) where $f : S \hookrightarrow T$ is an injection, and g is a d-coloring of the compliment of the image of f, i.e. a map from the compliment of the image of f to the set $[d] := \{1, \dots, d\}$.

If (f, g) and (f', g') are two composable morphisms, then we set

$$(f, g) \circ (f', g') = (f \circ f', h),$$

where

$$h(x) = \begin{cases} g(x) & \text{if } x \notin \mathrm{im} f \\ g'(f^{-1}(x)) & \text{otherwise.} \end{cases}$$

We observe that FI_1 is equivalent to the category FI of finite sets and injections. We also observe that FI_d has a fully faithful subcategory, whose objects are the sets $[n] = \{1, \dots, n\}$. We also refer to this subcategory as being FI_d.

Definition 2.11. Given a commutative ring k, an FI_d**-module over** k is a covariant functor $V : \mathrm{FI}_d \to \mathrm{Mod}_k$. Given an FI_d-module V, we will use V_n to denote $V([n])$.

If $(f, g) : [n] \to [m]$ is a morphism in FI_d, we write $(f, g)_*$ to denote the map $V(f, g)$. These maps are known as the **induced maps** of V. In the case where $n < m$, we refer to $(f, g)_*$ as a **transition map** of V.

Observe that the action of the FI_d-endomorphisms of $[n]$, for any n, make V_n into an $\mathfrak{S}_n$-representation over k. In this way, we may think of FI_d-modules as being sequences of $\mathfrak{S}_n$-representations, with n increasing, which are compatible with one another under the actions of the transition maps.

We will write FI_d-Mod to denote the category of FI_d-modules with natural transformations. Because the objects of FI_d-Mod are valued in an abelian category, it follows that FI_d-Mod is itself an abelian category.

Definition 2.12. Let m be a non-negative integer. The **free** FI_d**-module generated in degree** m, is defined on objects by

$$M(m)_n := k[\mathrm{Hom}_{\mathrm{FI}_d}([m], [n])],$$

the free k-module with basis indexed by the set $\mathrm{Hom}_{\mathrm{FI}_d}([m], [n])$. For each morphism (f, g), the induced map $(f, g)_*$ is defined on basis vectors by composition.

More generally, let W be a $k[\mathfrak{S}_m]$-module. Then we define the **free** FI_d**-module relative to** W, $M(W)$, as follows. For each positive integer n we set

$$M(W)_n := M(m)_n \otimes_{k[\mathfrak{S}_m]} W.$$

The induced maps are one again defined by composition on the first coordinate. Direct sums of free FI_d-modules of either type will also be referred to as being **free**.

Free modules are vitally important to the theory of FI_d-modules. In fact, their analogs appear to be fundamental objects in the representation theory of many other combinatorial categories. See [R] and [LY] for examples of this. One should note that there is a canonical isomorphism $M(k[\mathfrak{S}_m]) \cong M(m)$.

Definition 2.13. An FI_d-module V is said to be **generated in degree** $\leq m$ if there is a list of integers $\{m_i\}_{i\in I}$, with $m_i \leq m$ for all $i \in I$, and an exact sequence of FI_d-modules

$$0 \to K \to \bigoplus_{i\in I} M(m_i)^{n_i} \to V \to 0 \tag{1}$$

If the indexing set I can be taken to be finite, then we say that V is **finitely generated**.

Remark 2.14. Saying V is generated in degree $\leq m$ is equivalent to there existing a set $\{v_i\}_{i\in I} \subseteq \sqcup_{n=0}^{m} V_n$ such that no proper submodule of V contains every element of $\{v_i\}_{i\in I}$. This follows from the following important adjunction, for any $k[\mathfrak{S}_m]$-module W,

$$\mathrm{Hom}_{\mathrm{FI\text{-}Mod}}(M(W), V) = \mathrm{Hom}_{\mathfrak{S}_m}(W, V_m).$$

In particular, constructing the surjection in (1) is equivalent to choosing the set $\{v_i\}_{i\in I}$.

We write FI_d-mod to denote the category of finitely generated FI_d-modules. One of the main theorems about finitely generated FI_d-modules is the following.

Theorem 2.15 ([SS], Theorems 7.1.2 and 7.1.5)**.** *If V is a finitely generated FI_d-module over a Noetherian ring k, then all submodules of V are also finitely generated. Moreover, if k is a field, then there exists polynomials $p_1^V, \ldots, p_d^V \in \mathbb{Q}[x]$ such that*

$$\dim_k(V_n) = p_1^V(n) + p_2^V(n)2^n + \ldots + p_d^V(n)d^n$$

for all $n \gg 0$.

Prior to the provided source, the Noetherian property in the above theorem was proven for FI_d-modules over a field of characteristic 0 in [S, Theorem 2.3]. It was proven for FI-modules over a field of characteristic 0 in [CEF, Theorem 1.3]. Later [CEFN, Theorem A] proved this result for FI-modules over arbitrary Noetherian rings. The second part of the theorem, on dimensional stability, was proven in the case where k is a field of characteristic 0 in [S, Theorem 3.1]. It was proven for FI-modules over a field of characteristic 0 in [CEF, Theorem 1.5], and over an arbitrary field in [CEFN, Theorem B].

Definition 2.16. We call the function

$$n \mapsto \dim_k(V_n)$$

the **Hilbert function** of V. We say the Hilbert function of V is $o(d^n)$ if the polynomial $p_d^V \in \mathbb{Q}[x]$ from Theorem 2.15 is zero.

3. Representation Stability and FI_d-modules

In [CF], Church and Farb describe the phenomenon of representation stability. Following this, Church, Ellenberg, and Farb proved that representation stability could be equivalently stated in terms of finite generation of FI-modules [CEF, Theorem 1.13]. The goal of the first half of this section of the paper is to prove Theorem A, which suggests a kind of generalized representation stability. Following this we prove Theorem B.

We assume throughout the remainder of the paper that k is a field of characteristic 0.
t

3.1. Padded Partitions and Representations. t

Definition 3.1. Let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_h)$ be any partition, and let $(n_1, \ldots, n_r)$ be an r-tuple of integers such that $n_1 \geq n_2 \geq \ldots \geq n_r \geq |\lambda| + \lambda_1$. Then we define the **$r$-padded partition**

$$\lambda[n_1, \ldots, n_r] := (n_1 - |\lambda|, n_2 - |\lambda|, \ldots, n_r - |\lambda|, \lambda_1, \ldots, \lambda_h)$$

Similarly, we define the **r-padded representation**

$$S(\lambda)_{n_1,\ldots,n_r} := \begin{cases} S^{\lambda[n_1,\ldots,n_r]} & \text{if } n_1 \geq \ldots \geq n_r \geq |\lambda| + \lambda_1 \\ 0 & \text{otherwise.} \end{cases}$$

In this paper, we think of padding as a kind of parametrization. We have the following easy observation.

Lemma 3.2. *Let $\mu = (\mu_1, \dots, \mu_l)$ be a partition with at least r-rows. Then there exists a unique partition $\lambda = (\lambda_1, \dots, \lambda_h)$, as well as a unique collection of integers $n_1 \geq n_2 \geq \dots \geq n_r \geq |\lambda| + \lambda_1$, such that $\mu = \lambda[n_1, \dots, n_r]$.*

As a specific instance of the above lemma, if $\mu \vdash n$, then there is a unique partition λ such that $\mu = \lambda[n]$. This notation is convenient, as it allows us to uniformly describe the irreducible representations which appear in an FI_d-module. In particular, given an FI_d-module V, for all $n \geq 0$ we may write

$$V_n = V_n^{<d} \oplus \left(\bigoplus_{\lambda, n_1, \dots, n_d} c_{\lambda, n_1, \dots, n_d} S(\lambda)_{n_1, \dots, n_d} \right),$$

where we implicitly only allow $\lambda, n_1, n_2, \dots, n_d$ such that $|\lambda[n_1, \dots, n_d]| = n$, and the irreducible constituents of $V_n^{<d}$ are associated to partitions with strictly less than d rows. One may think of the above equation as expressing V_n as the sum of a grouping of all "small" length partitions, and a grouping of all "large" partitions. Theorem A will tell us that if V is finitely generated, then the multiplicities which appear in the above decomposition are stable in the appropriate sense.

Note that in the next section we will prove that if V is finitely generated, then the partitions λ which appear in the above decomposition have bounded size. This is analogous to the theorem from [CEF, Proposition 3.2.5], which states that the weight of a finitely generated FI-module is finite.

Definition 3.3. Let V be an FI_d-module and decompose V_n in the following way,

$$V_n = V_n^{<d} \oplus \left(\bigoplus_{\lambda, n_1, \dots, n_d} c_{\lambda, n_1, \dots, n_d} S(\lambda)_{n_1, \dots, n_d} \right),$$

where all the irreducible constituents of $V_n^{\leq d}$ correspond to partitions with strictly less than d rows. Then the **d-weight** $\mathrm{weight}^d(V)$ of V is the maximum value of $|\lambda|$ across all $S(\lambda)_{n_1, \dots, n_d}$ appearing in V_n for all $n \geq 0$. If no such integer exists, then we say $\mathrm{weight}^d(V) = \infty$.

3.2. **Free Modules in Characteristic 0.** We spend this section proving various lemmas and propositions about the FI_d-module $M(W)$. These will prove important to us in the sections which follow. We begin with a formula for $M(W)$ which is similar to that given in [CEF] for FI-modules.

Proposition 3.4. *Let W be a $k[\mathfrak{S}_m]$-module. Then for each n,*

$$M(W)_n = \bigoplus_{a=(a_1, \dots, a_d)} \mathrm{Ind}_{\mathfrak{S}_m \times \mathfrak{S}_a}^{\mathfrak{S}_n} W \boxtimes k,$$

where the sum is over compositions a of $n - m$ of length d.

Proof
Looking at the $\mathfrak{S}_n$-set $\mathrm{Hom}_{\mathrm{FI}_d}([m], [n])$, one quickly observes that the action of $\mathfrak{S}_n$ does not alter the colors which appear in a given morphism. In fact, it is easily checked that the orbits are in bijection with compositions of $n - m$ of length d. These orbits split the $\mathfrak{S}_n$-representation $M(m)_n$ into disjoint pieces. We will denote the piece associated to the composition $a := (a_1, \dots, a_d)$ by $M(m)_{n,a}$.

We may write

$$M(W)_n = M(m)_n \otimes_{k[\mathfrak{S}_m]} W = M(m)_{n,a} \otimes_{k[\mathfrak{S}_m]} W$$

where the sum is over the same set as in the statement of the proposition. We may consider the effect of tensoring W with $M(m)_{n,a}$ over $k[\mathfrak{S}_m]$ as equating elements (f, g) and (h, g') whenever f and h have the same image and $g = g'$. It follows that a basis for $M(W)_n$ is given by pairs $w_i \otimes (f, g)$, where $\{w_i\} \subseteq W$ is a basis for W, and (f, g) is such that f is monotone.

Fix a composition a of $n - m$, and consider the term $W \otimes_{k[\mathfrak{S}_m]} M(m)_{n,a}$. Let $(f, g) : [m] \to [n]$ be the pair of the standard inclusion, i.e. the inclusion which sends j to j, and the coloring which colors $m + 1, \dots, m + a_1$

the color 1, $m+a_1+1, \dots m+a_1+a_2$ the color 2, and so on. Then the subgroup $\mathfrak{S}_m \times \mathfrak{S}_a$ acts on the pure tensors $w \otimes (f,g)$ in the same way that it acts on $W \boxtimes k$. This gives us a map

$$\mathrm{Ind}_{\mathfrak{S}_m \times \mathfrak{S}_a}^{\mathfrak{S}_n} W \boxtimes k \to W \otimes_{k[\mathfrak{S}_m]} M(m)_{n,a}$$

It is easily seen that this map is surjective. On the other hand, the vector space on the left hand side has dimension

$$\dim(W)|\mathfrak{S}_n|/|\mathfrak{S}_m \times \mathfrak{S}_a| = \dim(W)\binom{n}{m, a_1, \dots, a_d}.$$

This is the same as the dimension of $W \otimes_{k[\mathfrak{S}_m]} M(m)_{n,a}$.

□

Observe that if $d = 1$, the above formula becomes

$$M(W)_n = \mathrm{Ind}_{\mathfrak{S}_m \times \mathfrak{S}_{n-m}}^{\mathfrak{S}_n} W \boxtimes k$$

This is the definition of $M(W)$ given in [CEF] and [CEFN].

The above description of the free module $M(W)$ clearly lends itself to applications of Pieri's rule. Indeed, using this simple combinatorial tool, we will be able to prove many representation theoretic properties of free modules.

Proposition 3.5. *Let $\lambda \vdash m$. Then* $\mathrm{weight}^d(M(S^\lambda)) = m$.

Proof
For the purposes of this proof, we write $M(\lambda) := M(S^\lambda)$. According to Proposition 3.4 we may write

$$M(\lambda)_n = \bigoplus_a \mathrm{Ind}_{\mathfrak{S}_m \times \mathfrak{S}_a}^{\mathfrak{S}_n} S^\lambda \boxtimes k$$

Thus, applying Pieri's rule, the irreducible factors S^μ which appear in $M(\lambda)_n$ will precisely be those for which there is a chain

$$\lambda \leq \lambda^{(1)} \leq \ldots \leq \lambda^{(d-1)} \leq \mu$$

such that $\lambda^{(i)}$ is obtained from $\lambda^{(i-1)}$ by adding some number (perhaps zero) of boxes to distinct columns. Assume that μ is a partition with at least d rows which is associated to an irreducible in $M(\lambda)_n$ and write $\mu = \nu[n_1, \dots, n_d]$. μ will be a constituent of $M(\lambda)_n$ if and only if one can obtain λ from μ by removing blocks from distinct columns in at most d steps. One can apply a "greedy" block removal algorithm which, in each step, removes all blocks which can be removed. It is clear that μ is a constituent of $M(\lambda)_n$ if and only if such an algorithm eventually yields λ. Observe that after the first step of the greedy algorithm applied to μ, one is left with the partition $(\mu_2, \dots, \mu_h)$. Therefore, after d steps one would be left with the partition ν. It follows that $|\nu| \leq |\lambda|$, and so $\mathrm{weight}^d(M(\lambda)) \leq |\lambda|$. Conversely, it is clear that the greedy algorithm applied to any d-padding of λ itself eventually terminates at λ precisely. This shows that $\mathrm{weight}^d(M(\lambda)) = |\lambda|$.

□

Observe that the above proposition implies that all finitely generated FI$_d$-modules have bounded d-weight.

Corollary 3.6. *If V is generated in degree $\leq m$, then* $\mathrm{weight}^d(V) \leq m$.

The proof of the above proposition also motivates the following observation about the irreducible constituents of free modules.

Proposition 3.7. *Let W be a finite dimensional $k[\mathfrak{S}_m]$-module, let μ be a partition, and let $c_{\mu,n_1,\dots,n_d}$ be the multiplicity of $S(\mu)_{n_1,\dots,n_d}$ in $M(W)$. Then the quantity $c_{\mu,n_1+l,\dots,n_d+l}$ is eventually independent of l.*

Proof
For a partition λ, once again write $M(\lambda) := M(S^\lambda)$. It suffices to prove the proposition for the module $M(\lambda)$, where $\lambda \vdash m$. In fact, we will show that the proposed stability happens when $l + n_d \geq \lambda_1 + |\lambda| = \lambda_1 + m$.

Let $n_d + l \geq \lambda_1 + m$, and assume that $S(\mu)_{n_1+l,\ldots,n_d+l}$ is a constituent of $M(\lambda)_n$. According to Pieri's rule, the multiplicity of $S(\mu)_{n_1+l,\ldots,n_d+l}$ in $M(\lambda)_n$ will be the number of ways to obtain $\mu[n_1+l,\ldots,n_d+l]$ from λ in at most d stages, such that blocks are added to distinct columns in each stage. In this proof, we visualize this process as follows. We begin with the tableaux for $\mu[n_1+l,\ldots,n_d+l]$, with a copy of the tableaux of λ in the upper left corner darkened. Note that if $\mu[n_1+l,\ldots,n_d+l]$ does not contain λ, then the multiplicity must be 0. If the relative sizes of μ and λ are such that λ is never contained in $\mu[n_1+l,\ldots,n_d+l]$, then our claim is proven. We may therefore assume that this is not the case. Note that, by our assumption on the size of l,

$$\lambda_1 \leq n_d + l - m \leq n_d + l - |\mu|,$$

from the previous lemma. It follows that a copy of λ will appear for our claimed value of l.

Beginning with our pre-darkened copy of λ inside the tableaux of $\mu[n_1+l,\ldots,n_d+l]$, we start darkening boxes in such a way that every column only receives one new darkened box and, when we are finished, the union of all darkened boxes is the tableaux of some partition. We repeat this process precisely d-times. Pieri's rule implies that the coefficient of $S(\mu)_{n_1+l,\ldots,n_d+l}$ in $M(\lambda)$ will be the number of ways that this procedure terminates with the entirety of $\mu[n_1+l,\ldots,n_d+l]$ darkened.

Because we only have d-steps, it follows that we must darken the first row of $\mu[n_1+l,\ldots,n_d+l]$ up to column $n_d+l-|\mu| = \mu[n_1+l,\ldots,n_d+l]_d$ in the first step. In fact, in the i-th step, row i must be filled in to at least column $\mu[n_1+l,\ldots,n_d+l]_d$. It follows that the only choices one has during the entire process are the rate in which the frame $(n_1-n_d,\ldots,n_{d-1}-n_d)$ is colored in, as well as the rate that the boxes below our pre-colored copy of λ are filled. The latter choices are entirely dependent on the relative sizes of λ and μ, and one does not gain or lose choices as l grows.

□

This is the major piece in proving Theorem A for free modules.

Proposition 3.8. *Let W be a finite dimensional $k[\mathfrak{S}_m]$-module, write $c_{\mu,n_1,\ldots,n_d}$ for the multiplicity of $S(\mu)_{n_1,\ldots,n_d}$ in $M(W)_n$, and let ϕ_n^i denote the induced map of the pair of the standard inclusion $[n] \hookrightarrow [n+1]$ with the color i. Then for all $n \geq m$ and all $l \gg 0$,*

1. *$\cap_i \ker(\phi_n^i) = 0$;*
2. *the images of the ϕ_n^i span V_{n+1} as an $k[\mathfrak{S}_{n+1}]$-module;*
3. *$c_{\mu,n_1+l,\ldots,n_d+l}$ is independent of l;*

Proof
As before, it suffices to prove the claim in the case of $M(\lambda) := M(S^\lambda)$. The first two properties above are clear from the definition of $M(\lambda)$. In fact, in this case all the maps ϕ_n^i are injective. Proposition 3.7 implies the third property.

□

3.3. **The Proof of Theorem A.** We begin with the following Theorem of Sam and Snowden.

Theorem 3.9 ([SS3])**.** *The Grothendieck group of finitely generated FI_d-modules over k is generated by classes of free objects, as well as classes of modules whose Hilbert function is $o(d^n)$.*

Recall from Theorem 2.15 that if V is a finitely generated FI_d-module, then there exists polynomials $p_1^V,\ldots,p_d^V \in \mathbb{Q}[x]$ such that $\dim_k(V_n) = p_1^V(n) + p_2V(n)2^n + \ldots + p_d^V(n)d^n$. One can interpret the above theorem of Sam and Snowden as saying that the $p_d(n)d^n$ term in this formula is realized in the Grothendieck group by sums of classes of free objects, while all lower terms arise from some slower growing modules.

Proof of Theorem A
First assume that V satisfies all conditions of Theorem A, and that V_n is finite dimensional for all n. Then there is some N such that the images of the ϕ_n^i span V_{n+1} as a $k[\mathfrak{S}_{n+1}]$-module for all $n \geq N$. For each $i \leq N$, let $\{v_{m,i}\}_{m=1}^{\dim_k(V_i)}$ be a basis for V_i. Then $\cup_{i,m}\{v_{m,i}\}$ is a generating set for V by definition. This

shows that V is finitely generated.

Conversely assume that V is finitely generated. The fact that V_n is finite dimensional for all n is clear from this. Next, set $K_n := \cap_i \ker(\phi_n^i)$. We make K_n into an $k[\mathfrak{S}_n]$-module by restricting the action of V_n, and we make the collection of K_n into an FI$_d$-module by making all the transition maps trivial. It follows from definition that K is a submodule of V. Theorem 2.15 implies that K must have a finite generating set $\{v_i\}$. Because the transition maps of K are all trivial, finite generation implies that it must be supported in only finitely many degrees. In particular, $K_n = \cap_i \ker(\phi_n^i) = 0$ for $n \gg 0$.

The definition of finite generation implies the second statement. It therefore remains to show that $c_{\lambda,n_1+l,\ldots,n_d+l}$ is eventually constant. Theorem 3.9 implies that it suffices to show the claim when $V = M(W)$, for some $k[\mathfrak{S}_m]$-module W, or when the Hilbert function of V is $o(d^n)$. Proposition 3.8 implies the statement when $V = M(W)$. We therefore assume that V has a Hilbert function which is $o(d^n)$. We will prove the stronger claim that $c_{\lambda,n_1+l,\ldots,n_d+l} = 0$ for $l \gg 0$. This is clear if $d = 1$, and so we assume that $d > 1$.

Our strategy to prove this stronger claim is to show that $\dim_k(S(\lambda)_{n_1+l,\ldots,n_d+l})$ grows, as a function of l, faster than our assumptions permit.

Let λ be a partition, fix $n_1 \geq \ldots \geq n_d \geq |\lambda| + \lambda_1$, and consider the tableaux associated to $\lambda[n_1+l, \ldots, n_d+l]$ for some positive integer l. Define $s := n_1 + l - |\lambda|$ and $r := \sum_i n_1 - n_i$. We note that r is the number of blocks missing to make the first d rows of $\lambda[n_1 + l, \ldots, n_d + l]$ into a rectangle, while s is the length of the first row. The hook formula tells us that

$$F(l) := \dim_k(S(\lambda)_{n_1+l,\ldots,n_d+l}) = \frac{(ds + |\lambda| - r)!}{\prod_{(i,j)\in\lambda[n_1+l,\ldots,n_d+l]} H(i,j)}.$$

Observe that we may multiply this expression by the polynomial $q(l) := \prod_{i=1}^{(d-1)|\lambda|+r}(ds + |\lambda| - r + i)$ to obtain

$$q(l)F(l) = \frac{(ds + d|\lambda|)!}{\prod_{i,j} H(i,j)}.$$

The function $q(l)$ is indeed a polynomial, as r is be independent of l.

At any box (i, j) in the first d rows of $\lambda[n_1 + l, \ldots, n_d + l]$ the hook with corner (i, j) is shorter horizontally than a hook which stretches to the final column s, and it's shorter vertically than if it stretched down past the remaining $d - i$ padded rows, and then through $|\lambda|$ blocks. This gives a lower bound,

$$q(l)F(l) \geq \frac{(ds + d|\lambda|)!}{(d-1+s+|\lambda|)(d-1+(s-1)+|\lambda|)\cdots(d-1+1+|\lambda|)(d-2+s+|\lambda|)\cdots(1+|\lambda|)C_\lambda}$$

where C_λ is a constant depending only on λ, which is the product of the sizes of hooks completely contained inside λ. Multiplying both sides of this inequality by terms of the form $(d - m + s + |\lambda|)$ with $1 \leq m \leq d$, the constant C_λ, and the constant $\frac{1}{\prod_{m=0}^{d-1}(d-m+|\lambda|)!}$, this becomes

$$\widetilde{q}(l)F(l) \geq \frac{(ds + d|\lambda|)!}{(s+|\lambda|)!(s+|\lambda|)!\cdots(s+|\lambda|)!} = \binom{ds + d|\lambda|}{s+|\lambda|, s+|\lambda|, \ldots, s+|\lambda|}$$

for some polynomial $\widetilde{q}(l)$.

Elementary combinatorics tells us that the multinomial coefficient on the right hand side of the above inequality is the largest among all coefficients with power $ds + d|\lambda|$ in d parts. This implies that it is bounded from below by the average of all such coefficients. The sum of all multinomial coefficients of a given power is an exponential, and the total number of such coefficients is a polynomial in the power. Thus,

$$\widehat{q}(l)F(l) \geq d^{ds+d|\lambda|}$$

for some polynomial $\widehat{q}(l)$. By assumption, our module has dimension growth $o(d^n)$, and s grows linearly with l. This shows that the multiplicity $c_{\lambda,n_1+l,\ldots,n_d+l}$ must eventually be zero, as desired.

□

3.4. **The Multiplicities of the Trivial Representation.** In this section we consider multiplicities of the trivial representations within in the module $V^{<d}$. In the next section, we will extend the main result of this section (Theorem 3.12) to multiplicities of more general irreducible representations (see Theorem B). We will find that the proof of the more general case amounts to a twist of the trivial representation case, and so we treat this case separately.

To better understand the multiplicity of the trivial representation, we make use of the coinvariants functor first introduced in [CEF] for FI-modules. The motivation for most of what follows can be found in [CEF, Section 3.1].

Definition 3.10. Let R denote the polynomial ring $k[x_1, \dots, x_d]$, and write Mod_R^{gr} for the category of non-negatively graded R-modules. The **coinvariants functor** $\Phi : \mathrm{FI}_d\text{-Mod} \to \mathrm{Mod}_R^{gr}$ is defined by

$$\Phi(V) = \bigoplus_n V_n \otimes_{k[\mathfrak{S}_n]} S^{(n)}.$$

where $S^{(n)} = k$ is the trivial representation.

One should think that taking coinvariants identifies transition maps whenever their coloring uses the same colors with the same frequency. After doing this, the maps ϕ_n^i, as in the statement of Theorem A, become the action by the variables x_i in R. One should also observe that if c_n is the multiplicity of the trivial representation in V_n, then $\dim_k \Phi(V)_n = c_n$.

We briefly record the following elementary properties of Φ.

Proposition 3.11. *The functor $\Phi : \mathrm{FI}_d\text{-Mod} \to Mod_R^{gr}$ enjoys the following properties.*

1. *Φ is exact;*
2. *$\Phi(M(m)) = R(-m)$, where $R(-m)$ denotes the graded twist of R by $-m$;*
3. *If V is finitely generated, then so is $\Phi(V)$.*

Proof
The first statement follows immediately from the fact that k is a field of characteristic 0.

Recall that for a composition $a = (a_1, \dots, a_d)$ of $n - m$ of length d, we write $M(m)_{n,a}$ to denote the submodule of $M(m)_n$ spanned by basis vectors (f, g), such that $|g^{-1}(i)| = a_i$. Each $M(m)_{n,a}$ is a permutation representation of $\mathfrak{S}_n$, where $\mathfrak{S}_n$ acts transitively. Therefore, $M(m)_n \otimes k$ has a basis in bijection with the set of all such compositions. The map ϕ_n^i acts on such a basis vector $(a_1, \dots, a_n)$ via $\phi_n^i(a_1, \dots, a_n) = (a_1, \dots, a_{i-1}, a_i + 1, a_{i+1}, \dots, a_n)$. There is therefore a natural isomorphism of graded R modules between $\Phi(M(m))$ and $R(-m)$ via

$$(a_1, \dots, a_d) \mapsto x_1^{a_1} \cdots x_n^{a_n}.$$

The final statement follows immediately from the first two, and the definition of finite generation.

□

It is a classically known fact that if M is a finitely generated graded R-module, then there is a polynomial $p \in \mathbb{Q}[x]$ such that $\dim_k M_n = p(n)$ for all $n \gg 0$ (See [E] for a proof). The discussion of this section therefore implies the following.

Theorem 3.12. *Let V be a finitely generated FI_d-module, and write c_n for the multiplicity of the trivial representation in V_n. Then there is a polynomial $p \in \mathbb{Q}[x]$ of degree $\leq d-1$ such that for all $n \gg 0$, $c_n = p(n)$.*

3.5. **The proof of Theorem B.** To finish the paper, we must prove Theorem B. This theorem was proven in the case wherein $\lambda = (n)$ in the previous section, and we will find that the proof of the general case follows from this. The main trick for the reduction of the general case to the previous case was shown to the author by Steven Sam. The author would like to send his thanks to Professor Sam for aiding him in this way.

We begin with some notation from the representation theory of categories.

Definition 3.13. Let $\mathcal{C}$ and $\mathcal{C}'$ be (small) categories. A **$\mathcal{C}$-module** over a ring R is a functor $V : \mathcal{C} \to R$-Mod. We say that a $\mathcal{C}$-module is **finitely generated** if there is a finite collection of elements $\{v_i\} \subseteq \bigsqcup_{X \in \mathcal{C}} V(\mathcal{C})$ which is contained in no proper submodule. Assume that V is a $\mathcal{C}$-module over some ring R, while V' is a $\mathcal{C}'$-module over the same ring. Then the **(exterior) tensor product of V an V'**, denoted $V \boxtimes V'$, is the $\mathcal{C} \times \mathcal{C}'$-module defined on points by

$$(V \boxtimes V')(X, X') = V(X) \otimes_R V'(X')$$

For the remainder of the section, We let V be a finitely generated FI_d-module over a field k of characteristic 0. We also fix a partition $\lambda \vdash m$, and set $c_{\lambda,n}$ to be the multiplicity of $S(\lambda)_n$ in V_n. The primary strategy for proving Theorem B is to twist a given FI_d-module with an appropriately chosen FI-module, so that the multiplicity of the trivial representation in the resulting module is equal to $c_{\lambda,n}$ at this point the previous section's results will imply Theorem B.

In [CEF, Proposition 3.4.1], Church, Ellenberg and Farb define a finitely generated FI-module $S(\lambda)$ with the property that

$$S(\lambda)_n = \begin{cases} 0 & \text{if } n < m \\ S^{\lambda[n]} & \text{otherwise.} \end{cases}$$

This is the FI-module with which we will twist V. To accomplish this formally, we first need some technical lemmas from the Gröbner theory of Sam and Snowden [SS].

Definition 3.14. Let $\mathcal{C}$ and $\mathcal{C}'$ be two categories, and let $\Psi : \mathcal{C} \to \mathcal{C}'$ be a (covariant) functor between these categories. Then we say that Ψ **satisfies property (F)** if given any object x of $\mathcal{C}'$, there is a finite list of objects $y_1, \ldots, y_r$ of $\mathcal{C}$ and morphisms $f_i : x \to \Psi(y_i)$, such that for any object y of $\mathcal{C}$ and any morphism $f : x \to \Psi(y)$ in $\mathcal{C}'$, there is a morphism $g : y_i \to y$ in $\mathcal{C}$ for some i such that $f = \Psi(g) \circ f_i$.

Property (**F**) was initially considered by Sam and Snowden in [SS]. In this work they show that this property is the key tool for connecting stability phenomena in the study of $\mathcal{C}'$-modules to stability phenomena in the study of $\mathcal{C}$-modules. For our purposes, we will only need the following lemmas.

Lemma 3.15 ([SS], Proposition 3.2.3)**.** *Let $\Psi : \mathcal{C} \to \mathcal{C}'$ be a functor, and write Ψ^* for the natural pullback $\Psi^* : \mathcal{C}'$-Mod $\to \mathcal{C}$-Mod. Then Ψ satisfies property (**F**) if and only if Ψ^* maps finitely generated $\mathcal{C}'$-modules to finitely generated $\mathcal{C}$-modules.*

Lemma 3.16. *Let $\Psi : \mathrm{FI}_d \to \mathrm{FI}_d \times \mathrm{FI}$ denote the functor defined by the assignments*

$$\Psi([n]) = ([n], [n]), \qquad \Psi((f, g)) = ((f, g), f).$$

*Then Ψ satisfies property (**F**).*

Proof
Let $([r], [s])$ be an object in $\mathrm{FI}_d \times \mathrm{FI}$, and assume that $r < s$. Looking at the definition of property (**F**), we set $y_i = [s]$ and allow the $f_i : ([r], [s]) \to \Psi(y_i)$ to vary across all of the (finitely many) morphisms $([r], [s]) \to ([s], [s]) = \Psi(y_i)$ in $\mathrm{FI}_d \times \mathrm{FI}$. Now given any $[l]$, and any morphism $f : ([r], [s]) \to \Psi([l]) \in \mathrm{FI}_d \times \mathrm{FI}$, we set $(h, h') : [s] \to [l] \in \mathrm{FI}_d$ to be the pair of the standard inclusion with a coloring chosen to agree with the coloring associated to f on the compliment of $[s]$ in $[l]$. It follows there is some map in $\mathrm{FI}_d \times \mathrm{FI}$, $g : ([r], [s]) \to ([s], [s])$ such that $\Psi(h, h') \circ g = f$. Such a g will be included among the f_i by how the f_i were chosen. The case where $r > s$ is the same.

□

This is all we need to prove Theorem B.

proof of Theorem B
Consider the $(\mathrm{FI}_d \times \mathrm{FI})$-module $V \boxtimes S(\lambda)$. This module is finitely generated, as both V and $S(\lambda)$ are. Lemmas 3.15 and 3.16 imply that $\Psi^*(V \boxtimes S(\lambda))$ is also a finitely generated FI_d-module. By construction, the multiplicity of the trivial representation in $\Psi^*(V \boxtimes S(\lambda))_n$ is precisely $c_{\lambda,n}$. Theorem 3.12 now implies Theorem B.

□

References

[CE] T. Church and J. S. Ellenberg, Homological properties of FI-modules and stability, `arXiv:1506.01022`.

[CF] T. Church and B. Farb, *Representation theory and homological stability*, Advances in Mathematics, (2013), 250-314.

[CEF] T. Church, J. S. Ellenberg and B. Farb, FI-*modules and stability for representations of symmetric groups*, Duke Math. J. 164, no. 9 (2015), 1833-1910.

[CEFN] T. Church, J. S. Ellenberg, B. Farb, and R. Nagpal, FI-*modules over Noetherian rings*, Geom. Topol. 18 (2014) 2951-2984.

[CST] T. Ceccherini-Silberstein, F. Scarabotti, and F. Tolli, *Representation Theory of the Symmetric Groups: The Okounkov-Vershik Approach, Character Formulas, and Partition Algebras*, Cambridge University Press, Cambridge, 2010.

[E] D. Eisenbud, *Commutative Algebra with a View Towards Algebraic Geometry*, Graduate Texts in Mathematics # 150, Springer-Verlag, New York, 1995.

[GL] W. Liang Gan, and L. Li, *Coinduction functor in representation stability theory*, J. London Math. Soc. (2015) 92 (3), 689-711.

[LY] L. Li, and N. Yu, *Filtrations and Homological Degrees of* FI-*modules*, `arXiv:1511.02977`.

[N] R. Nagpal, FI-*modules and the cohomology of modular representations of symmetric groups*, `arXiv:1505.04294`.

[R] E. Ramos, *Homological Invariants of* FI-*Modules and* FI_G-*Modules*, `arXiv:1511.03964`.

[S] A. Snowden, *Syzygies of Segre embeddings and* Δ-*modules*, Duke Math. J. 162 (2013), no. 2, 225-277.

[SS] S. Sam, and A Snowden, *Gröbner methods for representations of combinatorial categories*, `arXiv:1409.1670`.

[SS2] S. Sam, and A. Snowden, *GL-equivariant modules over polynomial rings in infinitely many variables*, Trans. Amer. Math. Soc. 368 (2016), 10971158.

[SS3] S. Sam, and A. Snowden, *GL-equivariant modules over polynomial rings in infinitely many variables. II*, in preparation.

[W] J. C. H. Wilson, FI_W-*modules and stability criteria for representations of the classical Weyl groups*, Journal of Algebra, Volume 420, 15 December 2014, Pages 269-332.

Department of Mathematics, University of Wisconsin - Madison.
E-mail address: `eramos@math.wisc.edu`